

DATA STREAM COMMUNICATION

CROSS REFERENCE

The present application claims priority under 35 U.S.C. §119 on US Provisional Application No. 60/472,648 filed on May 22, 2003.

FIELD OF THE INVENTION

The present invention is related to data communications, and to communication of a subset of data streams.

BACKGROUND OF THE INVENTION

Many data applications involve communications of a plurality of data streams. Audio and video conferencing, by way of example, may involve communication of audio data streams, video data streams, and other data streams. Within this exemplary application, some conferences simultaneously link many locations for communication of data such as video, audio, digital files, and the like. Each location may be communicating data streams from multiple cameras, microphones, and other sources. Under such circumstances, the bandwidth required for each individual location can be quite large. This large amount of required bandwidth can be costly, and may be unavailable for some conference attendees, thereby limiting their ability to participate.

One proposed solution is compression of signals. Even when using compressed signals, however, required bandwidth can be quite large. Also, compression may be accompanied by a loss of quality. Another proposed solution to this problem has been to send data streams from only one

conference attendee to sites that are operating using limited bandwidth. By way of particular example, one proposed method in the case of video conferencing using Real-time Protocol (RTP) data streams is the use of source specific multicast (SSM). This proposed solution leaves many problems unresolved, however. For example, SSM only allows for selection of data streams based on IP address, and is therefore not effective in the case where a single IP address is generating multiple streams since all of the streams with that address will be sent. Also, SSM can be complicated and costly to implement, with the result that it is not widely supported.

Unresolved problems in the art therefore remain.

SUMMARY OF THE INVENTION

The present invention is directed to methods for communicating at least one primary data stream. An exemplary method of the invention includes steps of monitoring a plurality of data streams communicated between a plurality of standard users, the data streams including a plurality of continuous data streams communicated from each of the plurality of standard users to others of the standard users. The method further includes steps of recognizing at least one primary data stream from the plurality of streams and communicating it to a primary user. In an exemplary application, each of the standard users is a virtual meeting attendee and is generating multiple video and audio data streams and that have relatively high bandwidth resources for communications. In the exemplary application the primary user is a meeting attendee with more limited bandwidth resources.

An additional embodiment of the invention is directed to a computer program product for communicating one or more primary data streams in a virtual meeting environment. The program product causes one or more computers to: communicate a plurality of continuous real-time data streams that include discretely packetized video and audio data between a plurality of standard users, and to identify a primary subset of the plurality of

continuous real-time data streams and communicate it to one or more primary users.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 is a schematic of a network;

FIG. 2 is a detailed schematic of individual standard users A-D of the network of FIG. 1;

FIG. 3 is a detailed schematic of an individual primary user 1 of the network of FIG. 1;

FIG. 4 is a detailed schematic of an individual primary user 2 of the network of FIG. 1; and,

FIG. 5 is a flowchart illustrating steps of one exemplary method of the invention.

DETAILED DESCRIPTION

Before discussing exemplary embodiments of the present invention in detail, it will be appreciated that the invention may be embodied in a method, a system, and/or in a computer program product. For example, a method of the invention may be carried out by one or more users using computers, and a program product of the invention may include computer executable instructions that when executed by a computer cause that computer to carry out a method of the invention. Further, a computer that contains a program product of the invention may embody a system of the invention. It will accordingly be appreciated that in describing a particular embodiment of the present invention, description of other embodiments may also be made. For example, it will be understood that when describing a method of the invention, a system and/or a program product of the invention may likewise be described.

Turning now to the drawings, FIG. 1 is a schematic of a network that is useful to describe an exemplary method of the invention. The network shown as a “cloud” 10 includes an interface 12 that links standard users A-D to one another. The term “interface” as used herein is intended to be broadly

interpreted as comprising one or more components for linking communications between users. It may include, for example, a computer having a plurality of communication ports, a bridge, a software component running on a computer that facilitates communications, a networking card, a modem, and the like. An exemplary interface 12 is a bridge with a plurality of ports. Those skilled in the art will appreciate that as used herein the term port is intended to be broadly interpreted as a physical or logical destination and/or origination point for communications. Examples of ports include but are not limited to, network cards, an IP address, a TCP or UDP port number, and the like. The network 10 may be a digital or analog communications network, with a packet switched protocol network being one example. It may be a land-based, physically wired network, or may be a wireless network. Also, the protocol between bridge 12 and the standard users A-D may be that of a server and clients.

In an exemplary application, the network 10 is useful to facilitate a virtual meeting between attendees that are physically present at each of the standard users A-D. As used herein the term “virtual meeting” is intended to be broadly interpreted as a sharing of real-time communications between participants that are not physically present with one another. Communications with each of the standard users A-D is carried out on a 2-way basis from the network 10, with data sent to and received from each of the standard users A-D over the communications lines represented as dashed line “pipes” 14. These may comprise physically wired connections, or may be wireless connections. Real-time video, audio, and other data may be sent from each of the standard users A-D to all others of the standard users A-D through the bridge 12 and over the communications lines 14. As used herein the term “real-time” is intended to broadly refer to a condition of generally corresponding to actual time. For example, data is real-time if it takes one minute of data playback to describe an event that took one minute to occur. The real-time data may be recorded and still be real-time.

Those knowledgeable in the art will appreciate that these communications may be carried out in any of a number of generally known

procedures. For example, known methods of one or more of uni-, multi-, or broad- cast may be used. Also, the data may be streaming. Each standard user site's cameras and microphones, by way of further example, may stream a continuous, real-time data stream on a particular multicast address and port number. As used herein the term continuous data stream is intended to broadly refer to a data stream sent in substantially continuous succession, although some degree of intermittency is contemplated. For example, a packetized data stream in internet protocol is "continuous."

One particular exemplary method for communicating and receiving the continuous data streams within the practice of the invention is according to the so-called "Real-time Transport Protocol" or "RTP." RTP is a widely supported Internet-standard protocol for the transport of real-time data, including audio and video. It can be used for media-on-demand as well as interactive services such as Internet telephony. RTP consists of a data and a control part. The latter is called RTCP. The data part of RTP is a thin protocol providing support for applications with real-time properties such as continuous media (e.g., audio and video), including timing reconstruction, loss detection, security and content identification.

Communications of the streaming real-time data between standard users A-D may be further appreciated through consideration of the schematic of FIG. 2 that shows standard user A in detail. This schematic may be considered to be exemplary of each of the other standard users B-D. The standard user A includes cameras 1, 2 and 3 that have been shown as elements 16, 18 and 20 respectively. These cameras may be trained on different people or things at the standard user A, with an example being camera 1 at a blackboard, camera 2 at a speaker, and camera 3 at an audience. Standard user A may also include a first microphone 22 for the speaker and a second microphone 24 for the audience. Although not illustrated, other cameras, microphones, computers, gateways, firewalls, multi-plexers, co/decoders and like devices may also be present.

In addition to sending audio and video data across the communication line 16, each of the standard users A-D also receives video, audio, and other data communicated from each of the other standard users A-D. Referring once again to the schematic of FIG. 2 by way of example, one or more projectors 26 may be provided to project real-time video images 28 from one or more of the other standard users B-D on a screen 30. Any number of video images may be provided that show video data in real-time from any number of other cameras or other sources located at standard users B-D. Further, the video images displayed may include charts, graphs, documents, other digital files, replayed video files, and the like. One or more speakers 32 may also be provided to play real-time audio from the other standard users B-D or other sources.

A particular example of a data file in addition to audio and video data includes shared documents having text, images, numerical values, and the like. For example, within a virtual meeting different users at different locations may desire to all work on a single document. In such circumstances, continuous updates of the document should be communicated between users.

A computer 34 may be provided to receive and send all of the video, audio, documents, digital files and other data at the standard user A. An application program, such as an RTP application, may be running on the computer 34 that provides signal coding/decoding, signal compression/decompression, coordinates receiving and sending of the data streams, and controls some other aspects of sending and receiving of the data streams. For example, the computer 34 may be used to control which or how many video images 28 are displayed on the screen 30, to size the images 28, to set audio levels for the speakers 32, and the like.

Each discrete data stream that is communicated has a unique identifier associated with it. By way of example, methods, program products, and systems of the invention may be practiced across packet switched networks 10 that are configured for carrying discretely packetized data communications, with internet protocol ("IP") communications being one example, and RTP

communications being a more specific example. In IP communications, continuous data is packed into discrete packets and provided with a destination address. The address may be a digital string, for instance, that identifies a port on the bridge 12. Each of the discrete packets of data may also include a unique identifier, such as a digital origination address.

The origination address may, for instance, be a digital string that identifies the computer 34 at standard user A. Within the RTP protocol, identifier information may be embedded into the header portion of individual packets by the RTP application programs running on the computers 34. For example, a particular data stream may have source identifying information such as an SSRC ("synchronization source" in RTP protocol) and/or another identifier that includes the user name, camera number, and IP address of the computer 34. The SSRC identifier carried in the RTP header and in various fields of RTCP packets is a random 32-bit number that is required to be globally unique within an RTP session.

According to the configurations of FIGS. 1-2, a real-time virtual meeting can occur between the standard users A-D. A rich, immersive, and extensive virtual meeting environment may be provided that includes video, audio, and other streaming data shared in real-time between multiple participants at multiple locations. For example, participants at each standard user A-D may simultaneously view and hear data from all others of the users A-D. Such meetings may be desirable for corporations, universities, government, and other organizations that have groups of people located remotely from one another that need to interact in a somewhat detailed manner.

It will be appreciated, however, that in conducting such virtual meetings, relatively large amounts of communication bandwidth may be required. Referring to the examples of FIGS. 1 and 2, each of the several cameras and microphones at each of the standard users A-D is sent as a streaming real-time data stream to each of the other standard users A-D. Table 1 summarizes the data communicated between the various primary users A-D:

	Outgoing Data Streams:	Incoming Data Streams:
Standard User A	A1 (CAM 1)	From STD. CLNT. B: B1-B5 (5 Streams)
	A2 (CAM 2)	From STD. CLNT. C: C1-C5 (5 Streams)
	A3 (CAM 3)	From STD. CLNT. D: D1- D5 (5 Streams)
	A4 (MIC 1)	
	A5 (MIC 2)	
Standard User B	B1 (CAM 1)	From STD. CLNT. A: A1-A5 (5 Streams)
	B2 (CAM 2)	From STD. CLNT. C: C1-C5 (5 Streams)
	B3 (CAM 3)	From STD. CLNT. D: D1-D5 (5 Streams)
	B4 (MIC 1)	
	B5 (MIC 2)	
Standard User C	C1 (CAM 1)	From STD. CLNT. A: A1-A5 (5 Streams)
	C2 (CAM 2)	From STD. CLNT. B: B1-B5 (5 Streams)
	C3 (CAM 3)	From STD. CLNT. D: D1-D5 (5 Streams)
	C4 (MIC 1)	
	C5 (MIC 2)	
Standard User D	D1 (CAM 1)	From STD. CLNT. A: A1-A5 (5 Streams)
	D2 (CAM 2)	From STD. CLNT. B: B1-B5 (5 Streams)
	D3 (CAM 3)	From STD. CLNT. C: C1-C5 (5 Streams)
	D4 (MIC 1)	
	D5 (MIC 2)	

As a result, in an exemplary meeting occurring according to the configuration of FIGS. 1-2 wherein each of the standard users A-D are meeting with one another, each standard user A-D is simultaneously streaming five discrete data streams (three cameras and two microphones) across the communication line 14 while also simultaneously receiving fifteen discrete

data streams (five from each of the three other standard users) across the line 14. The bandwidth required for each of the communications lines 14 is therefore substantial. In alternate configurations, all of the audio streams from each site may be bundled together. For example, standard user A data streams A4 and A5 (MIC 1 and 2) may be bundled into a single stream.

Through methods, systems, and program products of the invention, it is possible for additional meeting attendees to participate in virtual meetings without having the bandwidth capacity to support the communications of standard users A-D. Referring to FIG. 1 by way of example, two additional meeting attendees are primary user 1 and primary user 2. Each of these attendees communicates with the network 10 via a communications line 40 that is lower in bandwidth than the lines 14. These attendees may include, for example, a home-worker or a traveling salesperson that wishes to participate in the virtual meeting. Through methods, systems, and program products of the invention, the primary users 1 and 2 are able to participate in the virtual meeting with the standard users A-D in a seamless and effective manner.

The schematic of FIG. 3 shows primary user 1 in detail. Primary user 1 has a single camera 42, a single microphone 44, a monitor 46, and a speaker 48. The camera 42 and microphone 44 are useful, for instance, to allow an individual at primary user 1 to send video and audio data to the network 10 over the communications line 40. Accordingly, primary user 1 is set up for two-way participation in the virtual meeting. Primary user 2, on the other hand, is configured only for one-way (receive data streams only, no sending). FIG. 4 illustrates the primary user 2 that includes only a monitor 50 and a speaker 52.

Having now described a network configuration useful for practice of an exemplary method, program product, and system of the invention, the flowchart of FIG. 5 may be considered. Referring now to that figure as well as FIGS. 1-4, a method, system, and program product that provides for substantially seamless participation of the primary users 1 and 2 with the higher

bandwidth standard users A-D may be illustrated. Each of the standard users A-D communicates a plurality of continuous data streams including video, audio, and other data such as shared documents and the like to the bridge 12 (block 102). At the bridge 12, each of the continuous streams are communicated to all others of the standard users A-D (block 104). The continuous data streams communicated from the users A-D may include the various camera and microphone streams, and are preferably packetized. Additionally, each has a unique identifier that may include, for instance, data placed in the header portion of packets of the stream.

A primary selection command is also communicated from any of the standard users A-D to the bridge 12 (block 106). The primary selection command includes an identifier that identifies one or more of the plurality of data streams as a “primary” stream(s). The identifier may include information extracted from the header of the stream, such as an IP origination address, or a label such as “A1 (CAM 1)” for standard user A, camera 1. The primary selection command may include more than one identifier to identify more than one data stream. For example, one or both of standard user A camera 1 and microphone 1 might be identified as primary.

The primary selection command may originate from any of the standard users A-D, the primary users 1-2, or the bridge 12. A pre-determined selection may identify one or more primary streams. Additionally, a primary selection command may originate from any other location in communication with the network 10. For example, a meeting coordinator (not illustrated) may be monitoring or controlling the meeting, and may be located at some additional user (not illustrated) connected to the network 10. The command may be carried out, for example, by “highlighting” a particular data stream on one of the computers 34 using a mouse, a keyboard, or other selection tool. Pressing SHIFT-1, for example, may cause the identifier for a highlighted stream to be communicated in a primary selection command to a particular port at the bridge 12. Identifier information such as the SSRC (“synchronization

source” in RTP protocol) header information with the standard user name and camera number will be included in the primary selection command.

Once the primary selection command is received at the bridge 12, it is used to identify the primary data stream from the plurality of streams being communicated between standard users A-D (block 108). This may be accomplished through any of several procedures, with an example including steps of storing the primary stream identifier in a memory and comparing the identifier for each of the plurality of data streams being communicated between the standard users A-D to the primary identifier to identify the primary stream. It will be understood that as used herein the terms “storing” and “memory” are intended to be broadly interpreted. A wide variety of memories will be useful for practice with the invention, with examples including, but not limited to, static and dynamic memories, RAM, ROM, cached storage, virtual memory, and the like.

Once the primary stream has been identified, it is communicated to the primary users 1 and 2 (block 110). By way of example, if the primary command identified camera 1 and microphone 1 from standard user A as “primary,” these data streams would be communicated to the primary users 1 and 2 over the communications lines 40. The images from user A, camera 1 would then be played on the monitor 46 at primary user 1 and the audio from user A, microphone 1 to the speaker 48. Likewise, the speaker 52 and monitor 50 would play the audio and video, respectively, from standard user A, camera 1 and microphone 1.

In this manner, the primary users are able to receive a limited subset of the overall data streams being communicated between standard users A-D. By way of example, a single speaker may be seen and heard. Because of the limited data being communicated, substantially lower bandwidth is required.

Communication of the primary data streams to the primary users may be carried out in any number of particular manners. By way of example, the bridge 12 may monitor the traffic being communicated between the

standard users A-D, may replicate the primary stream(s) after recognizing them, and communicate them to primary ports on the that are in communication with the primary users 1 and 2. In an additional invention embodiment, the bridge 12 may be replaced by first and second interfaces in communication with one another. A first or standard interface may link the standard users over one or more ports, and the second or primary interface may link the primary users over one or more ports for communication of the primary data stream(s).

By way of a particular example, an application program running on the bridge 12 may join a multi-cast address across which the traffic between the standard users A-D is occurring, and will begin monitoring that traffic. Incoming packets that include RTP header information matching the selected primary stream RTP header information may be sent to primary users “listening” to bridge at a selected port number.

Variations of selection and communication of primary data streams are contemplated within practice of different embodiments of the invention. Protocols may be enforced for limiting the primary command use. For example, only particular users may be allowed to generate the primary command, or only particular data streams may be allowed to be identified as primary. Also, selection of a primary stream may automatically replace a previously selected primary stream (so that only one stream may be primary at a time). For example, a second primary identification command may cause a primary data stream identified in a first command to be replaced with a newly identified primary data stream. Streams may also be selected according to a priority, so that one stream may be selected as a “first” primary, one as a “second” primary, and so on. Some or all of these streams may be delivered to primary users in a prioritized order depending on the capacity, viewing preferences, or like factors of that user (e.g., all streams received if required capacity is present, only “first” stream if capacity will only support one stream, etc.).

The present invention also contemplates receiving primary selection commands from the primary users. For example, primary user 1 may be able to select a primary stream of a text list that is communicated to him showing which streams are available. Also, the primary users may be able to send a primary selection command that changes the primary stream. For example, primary user 1 may send a primary selection command that changes the primary stream he is receiving to the next ranked stream on a prioritized list. In still another example embodiment, the bridge 12 may present a compressed or summarized list of all of the streams available to each of the primary users, who may then select from that list which stream(s) should be designated as primary. The compressed or summarized list may be in the form of a text list of stream origination or description (e.g., "Std. user A, Camera 1), thumbnail image(s) of the stream(s) (e.g., a worldwide webpage), a random sampling of the streams sent in succession one by one, a sampling of one or more of the streams sent out in a non-continuous manner (e.g., updated only once every 1 sec or other period).

Primary users may be identified through a dialogue or other steps. For example, when initially registering with the bridge 12, users may be queried to determine available bandwidth. If a user specifies a certain bandwidth that is known to be below that required to receive all of the expected communication traffic, that user may be classified as a primary user and linked to the port(s) serving the primary users A-B. Or, steps of an automated discovery may be carried out whereby the available bandwidth for each user will be automatically determined. Users having bandwidth insufficient to carry the full set of communications will be classified as primary users. Additionally, querying or discovery steps may be carried out to determine what level of communications each of the primary users can support. Referring to FIGS. 1, 3 and 4 by way of example, through querying and/or discovery steps, it may be determined that primary user 1 has bandwidth capacity to support receiving two primary data streams and sending two primary streams, while primary user 2 only has resources sufficient to receive two streams. The results

of these discovery and querying steps may be used to establish rules for communication with the primary users (e.g., primary user 2 not allowed to send data streams, only to receive)

After making such a determination, primary users such as primary user 1 can provide streaming input data streams for bundling with the data streams being communicated between the standard users. Referring now to FIGS. 3 and 1, continuous data streams from the camera 42 and microphone 44 at primary user 1 may be communicated across the communications line 40 to the bridge 12. They may be received at the bridge 12 using one or more ports that are dedicated to primary user communications. These streams may then be bundled with all of the streams being communicated between the standard users A-D so that each of those standard users A-D will be able to see and hear an individual using the camera 42 and microphone 44.

In still an additional embodiment of the present invention, use of different primary streams is contemplated. For example, a different data stream could be selected as primary for each of primary users 1 and 2. Selection of primary streams may be used to determine which users get which primary streams. An additional aspect of primary stream selection may include a recipient identifier or other information that indicates which primary users will receive which primary streams. By way of particular example, a stream may be selected as primary for any users that log in to a meeting as a "salesman," while a second is primary for any who logs in as a "manager."

Embodiments of the present invention have been shown and described for illustration purposes in the environment of video/audio/data conferencing, and with particular examples of RTP format video and audio data. Those skilled in the art, however, will appreciate that the present invention will have benefits and value to many other applications. It is not limited to RTP, video, audio or any other particular data formats or types. It is believed, however, that the particular fields of video/audio/data conferencing and collaboration arts, and specifically those involving multiple participants

with multiple data streams, will realize great benefits from embodiments of the present invention.

Also, it will be appreciated that the present invention may be practiced in the form of a method, a system, or a computer program product; or in combinations of one or more of these forms. For example, a computer program product embodiment of the invention may include program instructions stored in one or more computer readable memories that when executed cause one or more computers to carry out steps of a method of the invention. It will be further appreciated that a computer program embodiment may be executed by more than one single computer located at more than one location across one or more networks. Referring to FIGS. 1-4 by way of example, it will be understood that some steps of the invention may occur at the bridge 12, some at the computers 34, 46, and 50, and some at one or more other computers in communication with the network 10. An additional invention embodiment may include a system that includes the one or more computers on which the program instructions reside.